arXiv:1608.06495v1 [cs.CV] 23 Aug 2016

Searching Action Proposals via Spatial Actionness Estimation and Temporal Path Inference and Tracking

Nannan Li[♯], *Dan Xu*[§], *Zhenqiang Ying*[♯], *Zhihao Li*[♯], *Ge Li*[♯]

[♯]Peking University Shenzhen Graduate School, Shenzhen, P.R.China
[§]DISI, University of Trento, Trento, Italy

Abstract. In this paper, we address the problem of searching action proposals in unconstrained video clips. Our approach starts from actionness estimation on frame-level bounding boxes, and then aggregates the bounding boxes belonging to the same actor across frames via linking, associating, tracking to generate spatial-temporal continuous action paths. To achieve the target, a novel actionness estimation method is firstly proposed by utilizing both human appearance and motion cues. Then, the association of the action paths is formulated as a maximum set coverage problem with the results of actionness estimation as a priori. To further promote the performance, we design an improved optimization objective for the problem and provide a greedy search algorithm to solve it. Finally, a tracking-by-detection scheme is designed to further refine the searched action paths. Extensive experiments on two challenging datasets, UCF-Sports and UCF-101, show that the proposed approach advances state-of-the-art proposal generation performance in terms of both accuracy and proposal quantity.

1 Introduction

Video action analysis is an important research topic for human activity understanding, and has gained a wide attention in recent years. A common task of video action analysis is action recognition, which aims to identify which type of action is occurring in a video volume [25,28,18]. Compared to action recognition, action detection is a more difficult task, as it requires not only determining the action class, but also localizing the action in the video. Similar to the object detection task, in which reliable object proposals play a crucial role in the detection performance [23], action proposal is also a fundamental problem in action detection.

This paper focuses on generating high-quality action proposals in both spatial compactness and temporal continuity. Existing works in the literature have made different efforts to address the problem, including segmentation-and-merging strategy [12,13,1,7], dense motion features [29,26,24], human-centric models [14,10], and object proposals based approaches [30,4]. Despite promising results achieved in these works, video action proposal generation is still a challenging problem

due to the complex spatio-temporal relationship modeling involved in the task. The problem can be considered as a task with two essential steps, namely spatial (*i.e.*frame-level) actionness estimation and temporal (*i.e.*video-level) action path generation. For one aspect, because of the large diversity and variation of human actions, it is difficult to generate robust frame-level actionness proposals which contain meaningful action motion patterns and are clearly discriminative from the background in unconstrained videos. For the other aspect, as a fact that the whole number of potential actionness regions on each frame usually has an exponential growth of the video duration [22], it is extremely impractical to calculate on all possible connections of the regions for generating the action paths and guaranteeing each of them associated with the same actor(s).

To tackle the above mentioned issues in the generation of action proposals, we propose a novel framework based on spatial actionness estimation from multiple cues and temporal action path extraction from a fast inference and tracking. Firstly, unlike previous works using selective search [23] or edge boxes [31] for generating object proposals for actionness estimation, we employ more action related cues including both human and motion. Secondly, a deep Faster-RCNN [15] network is trained and fine-tuned on augmented action detection datasets for obtaining accurate human proposals. Then action motion patterns with Gaussian Mixture Models are modeled for motion estimation of each human proposal. Both human and motion estimations are feed into a proposed forward and backward search algorithm for video-level action path generation. Finally, we use a tracking-by-detection approach to refine the action path by supplement actionness proposals missing in frames.

The key contributions of this paper include three folds: i) we construct an action detector at frame-level by taking both appearance and motion clues into account, which can handle the problem of detecting human with uncommon poses and discriminate actionness proposals containing meaningful motion patterns from the backgrounds; ii) We formulate the action path generation as a maximum convergence problem [5]. We propose an improved optimization objective for the problem and provide a greedy search algorithm to solve it. iii) Extensive experiments on UCF-Sports, UCF-101 datasets show that the proposed method achieves the state-of-the-art performance compared with other existing approaches.

2 Related Work

Traditionally, action localization or detection is performed by sliding window based approaches [20,11,3,27]. For instance, Siva *et al.*[20] proposed a supervised model based on multiple-instance-learning to slide over subvolumes both spatially and temporally for action detection. Instead of performing an exhaustive search through sliding over the whole video volumes, Oneata *et al.*[13] put forward a branch-and-bound search approach to achieve the time-efficiency. The main limitation of these sliding-window based approaches is that the detection

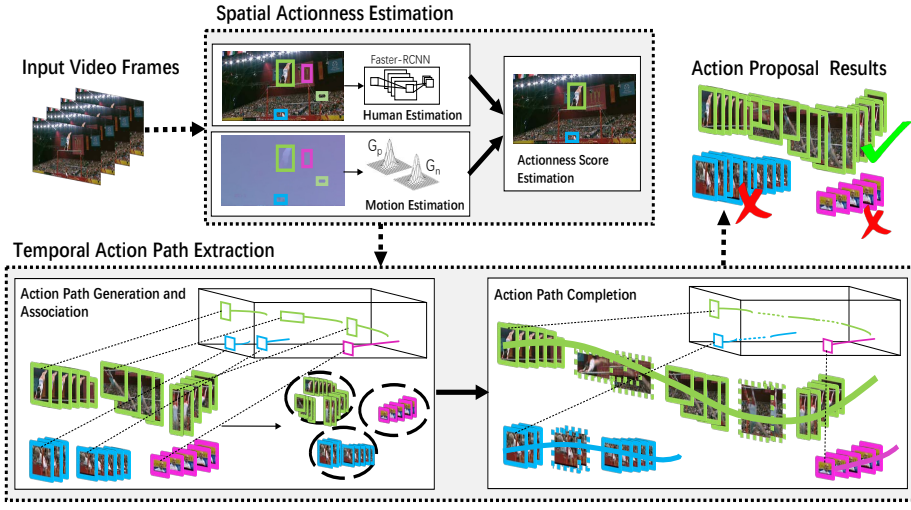

Fig. 1. The framework of the proposed action proposal generation approach.

results are confined by a video subvolume, and thus can not accurately capture the varying shape of the motion.

Some research works address the problem by employing segmentation-and-merging strategy. Generally, these methods include three steps: i) segment the video; ii) merge the segments to generate tube proposals; iii) represent tubes with dense motion features and construct action classifier for recognition. For instance, in [7] action tubes are generated by hierarchically merging super-voxels. However, accurate video segmentation is a difficult problem especially under unconstrained environments. To alleviate the difficulty encountered with segmentation, some other methods use a figure-centric based model. In [14] the human and objects are detected first and then their interactions are described. Kläser *et al.*[10] detect human on each frame and track the detection results across frames using optical flow. Our approach also utilizes tracking, via a more robust tracking-by-detection approach [6,9] based on a combined feature representation of color and shape.

Recently, some methods built upon generation of action proposals are presented. Gkioxari *et al.*[4] proposed to utilize Selective Search method for proposing actions on each frame, then scored those proposals by using features extracted by a two-streams Convolutional Neural Networks (CNN), and finally, linked them to form ation tubes. Philippe *et al.*[29] adopted the same feature extraction procedure, then utilized a tracking-by-detection approach to link frame-level detections, in combination with a class-specific detector. Our method replaces object proposal method and two-stream CNN with the Faster R-CNN model for calculation efficiency. The most related work to ours is that presented in [30], in which actionness score is calculated for each action path and then a greedy search method is used to generate proposals. Our work differentiates from theirs

in the following three aspects: i) we train a Faster R-CNN model for human estimation, which has a stronger ability to differentiate human from backgrounds; ii) compared with the optimization objective they proposed, our improved optimization objective simultaneously maximizing actionness score and member similarity in a path set, thus can effectively cluster the paths from the same actor into a group; iii) we utilize a tracking-by-detection approach to supplement the missing detections.

3 The Proposed Approach

The proposed approach takes video clip as input and generates action proposal results. The framework of our approach is illustrated in Fig. 1. The main procedure consists of two stages: spatial actionness estimation and temporal action path extraction. Firstly, bounding boxes at frame-level that may contain meaningful motion are extracted by simultaneously considering appearance and motion cues; then action paths corresponding to the same actor at video-level are generated and linked to obtain action proposals. The details of our method will be elaborated in the following sections.

3.1 Spatial Actionness Estimation

Human Estimation Detection of Human proposal is an important and heuristic step for action localization. We implement the human proposal detection employing the Faster R-CNN [15] pipeline with a VGG-16 model [19] pre-trained on ILSVRC dataset [17]. Faster R-CNN introduces a Region Proposal Network (RPN) that simultaneously predicts object bounding boxes and their corresponding objectness scores in near real-time speed. As the human detection task is a binary-classification problem, the output of the classification layer of Faster-RCNN network is revised to a two-way softmax classifier: one for the 'human' class and the other for the 'background' class. For action classes such as diving and swing, the appearance (especially for the shape and the pose) of the human changes significantly among the whole action duration. Therefore, the detection network fine-tuned on the standard PASCAL VOC 2007 dataset is unable to effectively detect the human under those circumstance. To handle the problem, we perform a data augmentation by merging the training data of the human class of PASCAL VOC 2007 and 2012, and rotating each training sample with seven different angles from $\frac{\pi}{4}$ to $\frac{7\pi}{4}$ with an interval of $\frac{\pi}{4}$. Let b_t^i denotes the bounding box for the i-th human proposal at t-th frame. The bounding box is represented as $[x, y, w, h]$, where w and h stand for width and height respectively, and (x, y) is the center. After training, for each bounding box b_*^* in the test video, a probability $S_h(b_*^*)$ can be estimated by the CNN network. By setting a probability threshold, human proposals with higher probability are kept for follow-up processing. A comparison of human detection results between original Faster R-CNN model and our refined one is showed in Fig. 2, from which it can be clearly observed that detection results from refined model are more precise and compact.

Fig. 2. Comparison of human detection results. The bounding boxes with red and green color are the ground truth and the detection results respectively. The 1-st and 3-nd columns are from the initial Faster R-CNN [15] (There is a missing detection in the 3-nd column); while the 2-nd and 4-th columns are from our fine-tuned model.

Motion Estimation Human cue provides important prior information for generating frame-level action proposals, however it is not sufficient to determine whether an action occurs, *e.g.*, human standing still. Thus we propose to further utilize motion cue for discarding false positive action proposals. The histograms of optical flow (HOF) [26] descriptor is used to describe the motion pattern of each human proposal. We construct two Gaussian Mixture Models (GMMs) $G_p(.)$ and $G_n(.)$ upon the HOFs, which represent the positive and negative proposal class respectively, to predict the probability of a motion pattern belonging to the actions or the background. HOFs calculated within bounding boxes of an Intersection-over-Union (IoU) overlapping with ground truth more than 0.5 are used as positive samples, while those with IoU overlapping less than 0.1 as negative samples. Given a test proposal b_t^i and its HOF h_i, we define the likelihood of b_t^i being a motion score using the predictions from two mixture of Gaussian models as:

$$S_m(b_t^i) = \sigma(G_p(h_i)/G_n(h_i)), \tag{1}$$

where $\sigma = 1/(1 + e^{-x})$ maps likelihood into the range [0, 1]. To reduce the influence induced by camera movement to optical flow calculation, we adopt the approach presented by [28] to estimate camera motion and subtract it to obtain robust optical flow.

Actionness Score Calculation The actionness score of a bounding box b_t^i consists of two parts: human detection score and motion score, and is defined as follows:

$$S(b_t^i) = S_h(b_t^i) + \lambda_p * S_m(b_t^i), \tag{2}$$

where λ_p is the parameter that balances the human estimation and motion estimation score.

3.2 Temporal Action Path Extraction

Problem Formulation Given action proposals on each frame, our goal is to find a set of action paths $\mathbf{P} = \{p_1, p_2, \ldots, p_i\}$, where $p_i = \{b_s^i, b_{s+1}^i, \ldots, b_e^i\}$ corresponds to a path that starts from s-th frame and ends at e-th frame. Yu and Yuan [30] formulate finding action path set $\mathbf{P}$ as a maximum set coverage

problem (MSCP) and propose an optimization objective maximizing actionness score. Inspired by their work, we formulate it as a MSCP with an improved optimization objective, which simultaneously maximizes actionness score and similarity among members within the path set $\mathbf{P}$. Formally, our optimization objective can be presented as follows:

$$\begin{aligned} &\max_{\mathbf{P}\in\Phi} \sum_{b_t\in\bigcup p_i} S(b_t) + \sum_{i,j} W(p_i,p_j) \\ &s.t. \quad |\mathbf{P}| \le N \\ &\qquad \mathbf{O}(p_i,p_j) \le \eta_p, \forall p_i,p_j \in \mathbf{P}, i \ne j, \end{aligned} \tag{3}$$

where $W(p_i,p_j)$ represents the similarity between action path p_i and p_j, and its definition will be explained in subsection action-path-association; $S(b_t)$ is the actionness score of bounding box b_t (cf. Eq. 2); Φ is action-path-candidate set; η_p is a threshold. The first constraint in Eq. 3 sets the maximum number of paths contained in $\mathbf{P}$; while the second constraint facilitates $\mathbf{P}$ to avoid generating redundant action paths that are overlapped. The overlapping of two paths is evaluated by $\mathbf{O}(p_i,p_j)$, which is defined as follows:

$$\mathbf{O}(p_i,p_j) = \frac{1}{\max(t_e^i,t_e^j) - \min(t_s^i,t_s^j)} \bullet \sum_{\max(t_s^i,t_s^j)\le t\le \min(t_e^i,t_e^j)} o(b_t^i,b_t^j). \tag{4}$$

In Eq. 4, $o(b_t^i,b_t^j)$ is defined as $\frac{\cap(b_t^i,b_t^j)}{\cup(b_t^i,b_t^j)}$, representing for IoU of two bounding boxs b_t^i and b_t^j.

Action Path Generation To solve the MSCP in Eq. 3, the action-path-candidate set Φ needs to be obtained first. We wish that Φ consists of spatio-temporal smooth path p_i whose consecutive elements b_t^i, b_{t+1}^i should satisfy the following two requirements:

$$\begin{aligned} o(b_t^i,b_{t+1}^i) &\ge \eta_o \\ \|C(b_t^i) - C(b_{t+1}^i)\| + \lambda_a \|H(b_t^i) - H(b_{t+1}^i)\| &\le \eta_f, \end{aligned} \tag{5}$$

where $o(b_t^i,b_{t+1}^i)$ represents IoU, as defined in Eq. 4; $C(b_t^i)$ and $H(b_t^i)$ stand for histograms of color (HOC) and histograms of gradient (HOG) of b_t^i, and λ_a is a trade-off balancing the weight of the two terms; η_o and η_f are thresholds. The first requirement in Eq. 5 ensures that consecutive bounding box b_t^i and b_{t+1}^i are spatially continuous; the second requirement ensures that b_t^i and b_{t+1}^i have similar appearance, thus the path p_i may follow the same actor.

To obtain Φ, we adopt the method proposed by [30] with minor modification to avoid generating much highly-overlap paths. The algorithm includes two stages: forward search and backward track. The aim of the former is to locate the

Algorithm 1 Forward Search and Backward Track

Input: bounding box score $S(b_t^i)$
Output: action path $p_k, k = 1, 2, \ldots, N$
1: $\tau_k = 0, b^k = \emptyset, k = 1, 2, \ldots, N$
2: **for** $t = 1 \rightarrow T$ **do**
3: **for** $i = 1 \rightarrow N_t^b$ **do**
4: $\tau(b_t^i) = \max_{b_{t-1}^j} \tau(b_{t-1}^j) + S(b_t^i)$
5: **end for**
6: step1: update each candidate (τ_k, b^k) with b_t^i that connects with b^k and has the largest score $\tau(b_t^i)$
7: step2: update (τ_N, b^N) as $(\tau(b_t^i,), b_t^i)$, if $\tau(b_t^i) > \tau_N$
8: **end for**
9: backward trace to locate $b_t^k, t = t_s, t_{s+1}, \ldots, t_e$ in p_k

end of a path; while that of the latter seeks to recover the whole path. The central idea is to maintain an updating pool of best Top-N path candidates, which is represented as $\Phi = (\tau_k, b^k), k = 1, 2, \ldots, N$, where τ_k is the score of path k and obtained by accumulating $S(b_t^k)$ of b_t^k it passes by; b^k is the bounding box of the end of k-th path. In the forward search, it also records an accumulated actionness score of each $b_t^i : \tau(b_t^i) = \max_{b_{t-1}^j} \tau(b_{t-1}^j) + S(b_t^i)$,where b_{t-1}^j and b_t^i satisfy the two requirements in Eq. 5. Given b_t^i at frame t, we update path candidate pool according to the following two steps: first, for each candidate $(\tau_k, b^k), k = 1, 2, \ldots, N$, if there exists any b_t^i connecting to b^k, then b^k will be replaced by b_t^i that has the largest $\tau(b_t^i)$; second, if the accumulated score of b_t^i is larger than the score of N-th proposal, i.e. $\tau(b_t^i) > \tau_N$, then (τ_N, b^N) is updated as $(\tau(b_t^i), b_t^i)$. After the forward search, a backward trace is performed to recover each b_t^k on the candidate path (τ_k, b^k). More specifically, for path p_k, we obtain $\{b_t^k : t = t_s, t_{s+1}, \ldots, t_e\}$ by solving the equation: $\tau_k = \sum_{t_s \leq t \leq t_e} S(b_t^k)$.

The pseudo-code of forward-backward search is illustrated in Algorithm 1. It takes bounding box score $S(b_t^i)$ as input data and outputs action paths $p_k, k = 1, 2, \ldots, N$. The lines 1 to 8 describe forward search and line 9 corresponds to backward trace. In line 3, N_t^b denotes the number of bounding box on frame t.

Action Path Association Once obtaining Φ, the MSCP in Eq. 3 can be solved. According to [5], the maximum set coverage problem is NP-hard but a greedy-search algorithm can achieve an approximation ratio of $1 - 1/e$. Here, we present a greedy-search solution to address the problem. In the beginning, we search for the candidate p_i with the largest action score τ_k in Φ, then add it into path set $\mathbf{P}$. Supposing that $\mathbf{P}$ has contained k action paths, we enumerate the rest paths in Φ and find the one that maximizes the flowing equation as the $k+1$-path p_i:

$$\arg\max_i \sum_{b \in p_i \cup p_1 \cup \ldots \cup p_k} S(b) + \sigma(1/k \cdot \sum_{j=1,2,\ldots,k} W(p_i, p_j)). \quad (6)$$

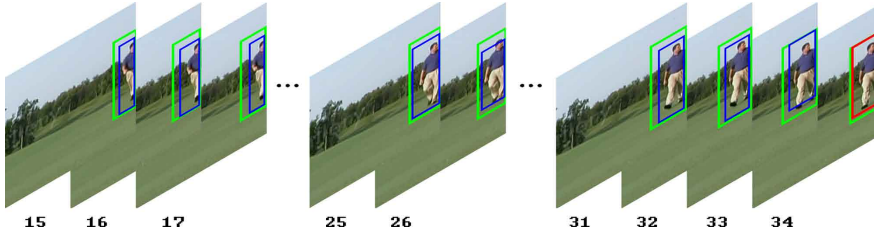

Fig. 3. Examples of tracking-by-detection results. The bounding boxes with green and red color are the groundtruth and the detected frame-level human bounding boxes respectively, and those with blue color are obtained by our tracking-by-detection strategy. All the missing human targets before frame 34 are perfectly located by the tracking approach.

In Eq. 6, $W(p_i, p_j)$ represents the similarity of action path p_i and p_j, and is defined as: $W(p_i, p_j) = 1/(\|C(p_i) - C(p_j)\| + \lambda_a \|H(p_i) - H(p_j)\|)$, where $C(p_*)$ and $H(p_*)$ represent the cluster centers of HOC and HOG of bounding boxes from path p_* respectively. The larger value of $W(p_i, p_j)$, the more likely that the paths p_i and p_j follow the same actor. To reduce redundant paths in set $\mathbf{P}$, the newly added path p_i should satisfy the constraint in Eq. 4.

Action Path Completion As human detection may miss hitting in some frames, the track obtained by connecting the paths in $\mathbf{P}$ will have temporal gaps. To get a temporal-spatial continuous track of an actor, we fill the gaps by using tracking-by-detection approaches [29]. We train a linear SVM as frame-level detector. The initial set of positives consist of bounding boxes in set $\mathbf{P}$, while negatives compose of bounding boxes excluded from set $\mathbf{P}$ and boxes that are randomly selected around positives with the IoU less than 0.3. Given the detection region b_t on frame t, we intend to find the most likely location on frame $t+1$ where the human detection is missed. Firstly, we map b_t to b'_{t+1} with the shift of the median of optical-flow inside region b_t; secondly, construct a search region $\overline{b'_{t+1}}$ by extending the height and width of b'_{t+1} to half past one times of original length; thirdly, scan $\overline{b'_{t+1}}$ with a set of windows whose ratio between width and length varies in a range $[0.8, 1.2]$ to adapt possible size change of an actor. The best region b_{t+1} is selected as the one that maximizes the following equation:

$$b_{t+1} = \arg \max_{\gamma \in N(\overline{b'_{t+1}})} S_f(\gamma), \tag{7}$$

where $N(\overline{b'_{t+1}})$ represents the window set produced by scanning $\overline{b'_{t+1}}$ and $S_f(\cdot)$ is the SVM detector whose input feature is chosen as the combination of HOC and HOG. After obtaining b_{t+1}, we update the SVM detector by adding b_{t+1} as a positive sample and boxes around b_{t+1} with the IoU less than 0.3 as negatives. An example of how the tracking approach supplementing missing detections is illustrated in Fig. 3.

3.3 Action Proposal Generation

The spatio-temporal continuous track can be considered as an action tube that focuses on an actor from appearing until disappearing. For each action tube, if its duration is larger than a specified threshold (*e.g.* 20), we regard it as an action proposal, denoted as $\mathscr{T}$.

4 Experiment

In this section, we describe the details of the experimental evaluation of the proposed approach, including datasets and evaluation metricts, implementation details, an analysis of the proposed approach and the overall performance comparison with state of the art methods.

4.1 Datasets and Evaluation Metric

We evaluate the performance of the proposed action proposal approach on two publicly available action-detection datasets: UCF-Sports [16] and UCF-101 [21].

UCF-Sports UCF-Sports dataset consists of 150 short videos of sports collected from 10 action classes. It has been widely used for action localization. The videos are truncated to contain a single action and bounding box annotation is provided for each frame.

UCF-101 UCF-101 dataset has more than 13000 videos that belong to 101 classes. In a subset of 24 categories, human actions are annotated both spatially and temporally. Compared with UCF-Sports, only a part of videos (74.6%) are trimmed to fit the motion.

Evaluation Metric To evaluate the quality of the action proposal $\mathscr{T}$, we follow the metric proposed by [24]. More specifically, the estimation is based on the mean IoU value between action proposal $\mathscr{T}$ and ground truth $\mathbf{G}$, which is defined as: $IoU(\mathbf{G}, \mathscr{T}) = \frac{1}{|\mathbf{C}|} \sum_{t \in \mathbf{C}} o(\mathbf{G}_t, \mathscr{T}_t)$, where $\mathbf{G}_t$ and $\mathscr{T}_t$ are the detection bounding box and ground truth on t-*th* frame respectively; $o(.,.)$ is the IoU value that is defined in Eq. 4; $|\mathbf{C}|$ is the set of frames where either the detection result or the ground truth is not null. An action proposal is considered as true-positive if $IoU(\mathbf{G}, \mathscr{T}) \geq \eta$, where η is a specified threshold. In the following passage, η is set as 0.5 if not specified.

4.2 Implementation Details

The human estimation is implemented under the Caffe platform [8] and based on the Faster R-CNN pipeline with a VGG16 model for parameter initialization as described in Sec. 3.1. We use a four-step alternating training strategy [15] to optimize two pipelines (*i.e.*RPN and Fast RCNN) of the whole network. For training the RPN pipeline, the same settings of scales and aspect ratios are used as in [15]. For training the Fast RCNN pipeline, the mini-batch size is set to 128, and the ratio of positive to negative samples is set to 1:4. The network is

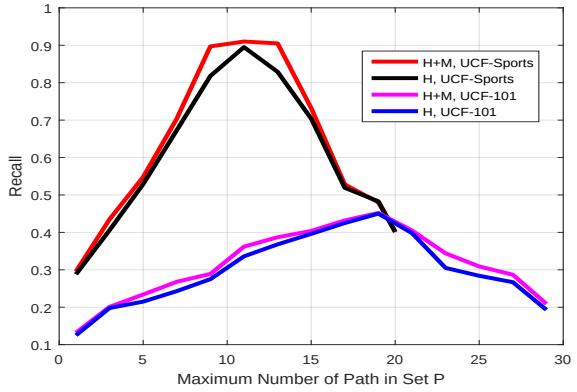

Fig. 4. Recall vs. maximum number of path in set $\mathbf{P}$ under different test settings.

trained with Stochastic Gradient Descent (SGD) with an initial learning rate of 0.001 and drop by 10 times at every the 5-th epochs, and the momentum and weight decay are set as 0.9 and 0.0005 respectively.

For the motion estimation of the bounding boxes, the number of components of GMMs is set to the same as the number of action categories. For constructing GMMs, we use randomly selected 1/3 of the video clips in UCF-101 for training and test on UCF-Sports dataset. While test on UCF-101 dataset, all the video clips of UCF-Sports are used for training. This setting is for a fair comparison with existing non-learning based methods which test on the whole dataset. The number of action paths in a candidate set Φ is set to 50 for UCF-Sports and 100 for UCF-101, as the latter one has a longer duration of action videos on average, and hence may contains more action-path segments. The value of N in Eq. 3 (*i.e.*the maximum number of paths in set $\mathbf{P}$) is set as 12 for UCF-Sports and 18 for UCF-101. For each video clip, we propose at least one path set $\mathbf{P}$, while a path set $\mathbf{P}$ is generated, the paths $\{p_i, i = 1, 2, \ldots, N\}$ in $\mathbf{P}$ are removed from the candidate set Φ and the greedy search algorithm is performed again to find a new path set $\mathbf{P}^{'}$ until the duration of the longest path $p_i^{'}$ ($p_i^{'} \in \mathbf{P}^{'}$) is less than 10.

4.3 Analysis of the proposed approach

We analyze the performance of the proposed approach from different aspects, including the sensitivity of the parameter N (*i.e.*, the maximum number of path in set $\mathbf{P}$), the influence of actionness estimation based on human appearance and motion cues, the number of generated proposals and the runtime analysis.

Fig. 4 shows the recall performance of our approach using different actionness estimation schemes by varying the value of N. From Fig. 4, it can be observed that the proposed approach achieves the best performance when the value of N is in the range $[9, 14]$ on UCF-Sports, and the performance degrades significantly when N is far from this range. The optimal value of N for UCF-101 is larger

than that for UCF-Sports. The reason is probably that the video clips of UCF-101 have longer duration than UCF-Sports on average, and thus the action path is more likely to be separated into multiple segments.

We can also observe that the proposed approach using both the human appearance and motion cues for actionness estimation (*i.e.*H + M) yields better recall performance than that using only the human appearance cue (*i.e.*H) on the two datasets. This demonstrates our initial intuition that employing multiple action-related cues for actionness estimation can help to further improve the performance of action proposal generation.

As also shown in Fig. 4, at the best performance point of recall on UCF-Sports, the number of the generated action proposals of our approach is only 13, and it is significantly less compared with state-of-the-art methods on the same recall performance level (see Table 1). The notable improvement is mainly due to the precise human estimation from our fine-tuned Faster-RCNN model, and the modified forward-backward search algorithm for generating candidate set Φ. Compared to [30], the improved optimization objective leverages appearance similarity among paths for effectively separating different actors. Fig. 5 illustrates the improvement by our approach. It can be observed that the action path generated by the proposed approach is correctly associated to the same actor. More examples of the action-proposal generation results on UCF-Sports and UCF-101 are shown in Fig. 8 and Fig. 9, respectively.

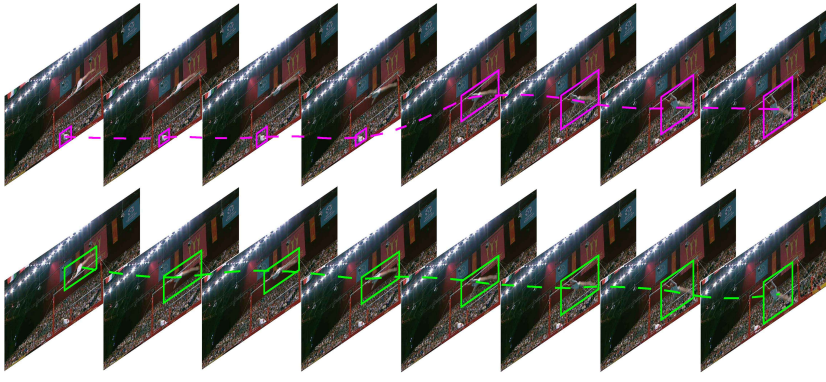

Fig. 5. Examples of action-path generation results. The 1-st row shows the results obtained from [30] (The action path contains an irrelevant actor within the first few frames); the 2-nd row is our results, where the main actor is correctly tracked.

The runtime of the proposed approach includes three parts: (i) spatial actionness estimation: Faster RCNN for human estimation takes around 0.1 seconds per frame (s/f) and GMM-HOF for motion estimation takes around 1 s/f; (ii) temporal action path extraction: the average runtime of this step is 0.09 s/f; and (iii) action path completion takes 0.5 s/f. In summary, the average runtime of the approach is 1.69 s/f. We conduct the runtime analysis on the UCF-Sports dataset

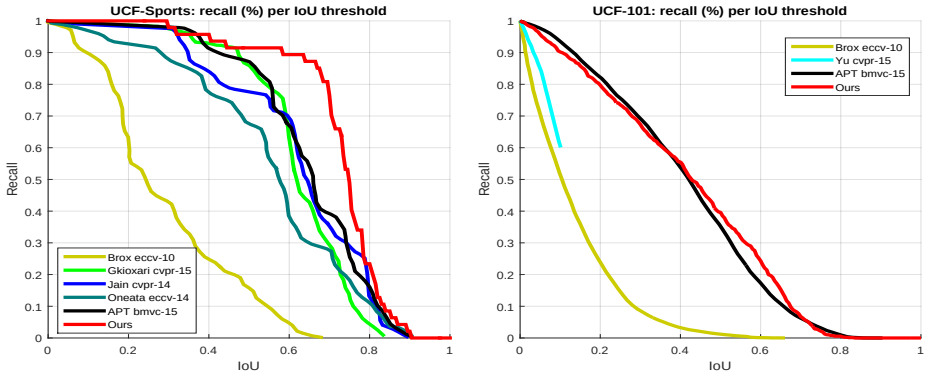

Fig. 6. Recall vs. IoUs on UCF-Sports and UCF-101 datasets.

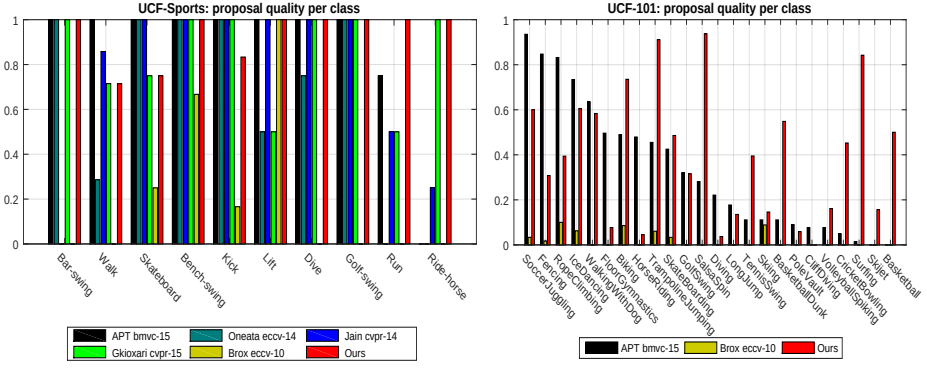

Fig. 7. Recall performance on each action category on UCF-Sports and UCF-101 datasets.

with an image resolution of 720 x 404, and based on hardware configurations of an Nvidia Tesla-K80 GPU, 3.4 GHz CPU and 4 GB memory.

4.4 Overall Performance

We compare the performance of the action proposal generation of the proposed approach with state-of-the-art methods on UCF-Sports and UCF-101. We vary the value of IoU (η) in [0, 1], and plot recall as a function of η. Fig. 6 shows Recall vs. IoU curves of difference approaches. It is clear that our approach obtains a significant performance gap over the state-of-the-art methods on UCF-Sports (The recall of ours is above 0.7 when η is even at 0.7, while the others are below 0.4.), and achieves very competitive performance on UCF-101. Fig. 7 also shows the recall performance for each action category. We can observe that our approach presents superior recall performance on almost all action categories except for few classes (*e.g.*, Walk and Kick) on UCF-Sports, and on UCF-101, ours

Table 1. Quantitative performance comparison of the action proposal generation with state-of-the-art methods with commonly used metrics.

	ABO	MABO	Recall	#Proposals
UCFSprots				
Brox & Malik, ECCV 2010 [2]	29.84	30.90	17.02	4
Jain *et al.*, CVPR 2014 [7]	63.41	62.71	78.72	1,642
Oneata *et al.*, ECCV 2014 [13]	56.49	55.58	68.09	3,000
Gkioxari & Malik, CVPR 2015 [4]	63.07	62.09	87.23	100
APT, BMVC 2015 [24]	65.73	64.21	89.36	1,449
Ours	**89.64**	**74.19**	**91.49**	12
UCF 101				
Brox & Malik, ECCV 2010 [2]	13.28	12.82	1.40	3
APT, BMVC 2015 [24]	40.77	39.97	35.45	2,299
Ours	**63.76**	**40.84**	**39.64**	18

greatly outperforms the comparison methods on action classes such as Biking, Surfing and TrampolingJumping.

From Fig.7, it can be also noticed that the performance on UCF-101 is inferior to that on UCF-Sports. The reason is probably because the testing action video clips of UCF-101 have more dynamically and continuously varying size of actors and are more untrimmed than UCF-Sports. Finally, we report the overall performance on the two datasets in Table 1 using several commonly used metrics, including ABO (Average Best Overlap), MABO (Mean ABO over all classes) and the average number of proposals per video. The results further confirm the superior action proposal generation performance achieved by our approach. It should be noted that on the same level recall performance, our approach generates relatively much smaller number of action proposals for each video clip than the other methods, which is especially important for reducing the computational complexity for the follow-up applications such as action recognition and action interaction modeling.

5 Conclusions

A novel framework for action proposal generation in video has been presented in this paper. Given an unconstrained video clip as input, it generates spatial-temporal continuous action paths. The proposed approach is built upon actionness estimation leveraging both human appearance and motion cues on frame-level bounding boxes, which are produced by a Faster-RCNN network trained on augmented datasets. Then we search spatial-temporal action paths via linking, associating, and tracking the bounding boxes across frames. We formulate the association of action paths belonging to the same actor as a maximum set coverage

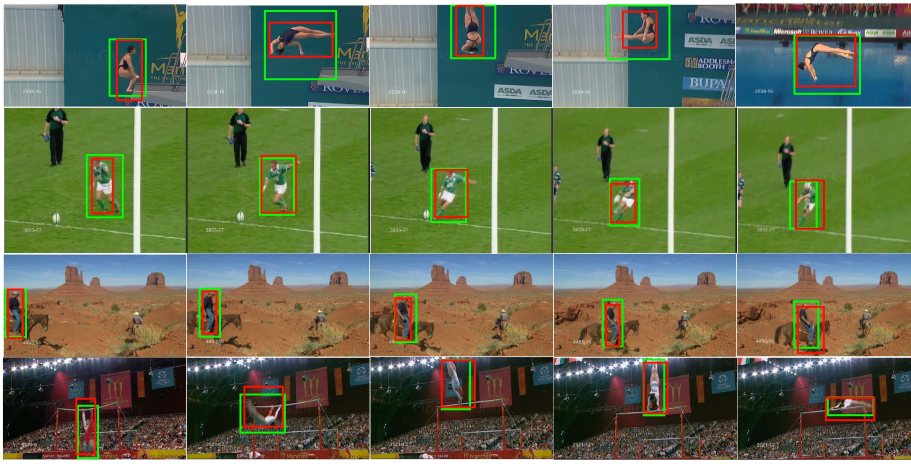

Fig. 8. Examples of action-proposal generation results on UCF-Sports. The bounding boxes with green and red color are the ground truth and the action proposal, respectively.

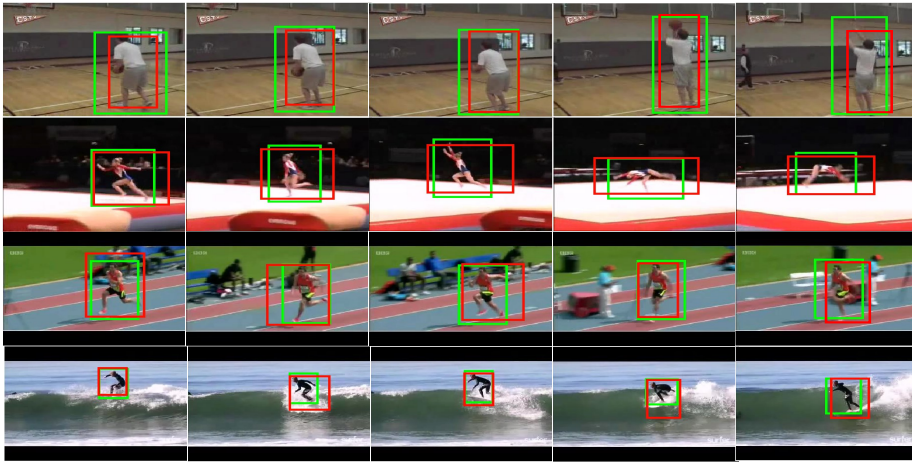

Fig. 9. Examples of action-proposal generation results on UCF-101. The bounding boxes with green and red color are the ground truth and the action proposal, respectively.

problem and propose a greedy search algorithm to solve it. Experiments on two challenging datasets demonstrate that our approach produces more accurate action proposals with remarkably less proposals compared with the state-of-the-art approaches. Based on the observation on the experimental results, the proposed approach is especially effective when the action video clip contains only one actor. In the future, we will explore using CNN networks for better actionness

estimation of video frames, and consider recurrent neural networks for modeling the action paths for action recognition.

References

1. Bergh, M., Roig, G., Boix, X., Manen, S., Gool, L.: Online video seeds for temporal window objectness. In: ICCV (2013)
2. Brox, T., Malik, J.: Object segmentation by long term analysis of point trajectories. In: ECCV (2015)
3. Gaidon, A., Harchaoui, Z., Schmid, C.: Temporal localization of actions with actoms. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 35(11), 2782–2795 (2013)
4. Gkioxari, G., Malik, J.: Finding action tubes. In: CVPR (2015)
5. G.L. Nemhauser, L.A.Wolsey, M.F.: An analysis of approximations for maximizing submodular set functions i. Mathematical Programming 14, 265–294 (1978)
6. Hare, S., Saffari, A., Torr, P.H.: Struck: Structured output tracking with kernels. In: ICCV (2011)
7. Jain, M., Gemert, J., Jégou, H., Bouthemy, P., Snoek, C.: Action localization with tubelets from motion. In: CVPR (2014)
8. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: Convolutional architecture for fast feature embedding. arXiv preprint arXiv:1408.5093 (2014)
9. Kalal, Z., Mikolajczyk, K., Matas, J.: Tracking-learning-detection. IEEE Transactions on Pattern Analysis and Machine Intelligence(PAMI) 34(7), 1409–1422 (2012)
10. Kläser, A., Marszałek, M., Schmid, C., Zisserman, A.: Human focused action localization in video. In: Trends and Topics in Computer Vision, pp. 219–233. Springer (2010)
11. Laptev, I., Pérez, P.: Retrieving actions in movies. In: ICCV (2007)
12. Ma, S., Zhang, J., Ikizler-Cinbis, N., Sclaroff, S.: Action recognition and localization by hierarchical space-time segments. In: ICCV. pp. 2744–2751 (2013)
13. Oneata, D., Revaud, J., Verbeek, J., Schmid, C.: Spatio-temporal object detection proposals. In: ECCV (2014)
14. Prest, A., Ferrari, V., Schmid, C.: Explicit modeling of human-object interactions in realistic videos. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 35(4), 835–848 (2013)
15. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NIPS (2015)
16. Rodriguez M D, Ahmed J, S.M.: Action mach a spatio-temporal maximum average correlation height filter for action recognition. In: CVPR (2008)
17. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. International Journal of Computer Vision(IJCV) 115(3), 211–252 (2015)
18. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: NIPS (2014)
19. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)

20. Siva, P., Xiang, T.: Weakly supervised action detection. In: BMVC. vol. 2, p. 6 (2011)
21. Soomro K, Zamir A R, S.M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. In: CoRR (2012)
22. Tran, D., Yuan, J., Forsyth, D.: Video event detection: From subvolume localization to spatiotemporal path search. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 36(2), 404–416 (2014)
23. Uijlings, J.R., van de Sande, K.E., Gevers, T., Smeulders, A.W.: Selective search for object recognition. International Journal of computer vision (IJCV) 104(2), 154–171 (2013)
24. Van Gemert, J.C., Jain, M., Gati, E., Snoek, C.G.: Apt: Action localization proposals from dense trajectories. In: BMVC (2015)
25. Wang, H., Kläser, A., Schmid, C., Liu, C.L.: Action recognition by dense trajectories. In: CVPR (2011)
26. Wang, H., Kläser, A., Schmid, C., Liu, C.L.: Dense trajectories and motion boundary descriptors for action recognition. International Journal of computer vision (IJCV) 103(1), 60–79 (2013)
27. Wang, H., Oneata, D., Verbeek, J., Schmid, C.: A robust and efficient video representation for action recognition. International Journal of Computer Vision (IJCV) pp. 1–20 (2015)
28. Wang, H., Schmid, C.: Action recognition with improved trajectories. In: ICCV (2013)
29. Weinzaepfel, P., Harchaoui, Z., Schmid, C.: Learning to track for spatio-temporal action localization. In: ICCV (2015)
30. Yu, G., Yuan, J.: Fast action proposals for human action detection and search. In: CVPR (2015)
31. Zitnick, C.L., Dollár, P.: Edge boxes: Locating object proposals from edges. In: ECCV (2014)